MASSACHUSETTS
INSTITUTE OF
TECHNOLOGY

Fault-tolerant Computing

6.033 Lecture 15

Hari Balakrishnan

Spring 2015

Your computer restarted because of a problem. Press a key or wa
seconds to continue starting up.

Votre ordinateur a redémarré en raison d'un problème. Pour pours
e redémarrage, appuyez sur une touche ou patientez quelques se

El ordenador se ha reiniciado debido a un problema. Para continu
el arranque, pulse cualquier tecla o espere unos segundos.

Ihr Computer wurde aufgrund eines Problems neu gestartet. Drüc
Sie zum Fortfahren eine Taste oder warten Sie einige Sekunden.

問題が起きたためコンピュータを再起動しました。このまま起動する場合
いずれかのキーを押すか、数秒間そのままお待ちください。

电脑因出现问题而重新启动。请按一下按键，或等几秒钟以继续


Windows
A fatal exception 0E has occurred at 0028:C00068F8 in PPT.EXE<01> +
000059F8. The current application will be terminated.
* Press any key to terminate the application.
* Press CTRL+ALT+DEL to restart your computer. You will
lose any unsaved information in all applications.
Press any key to continue

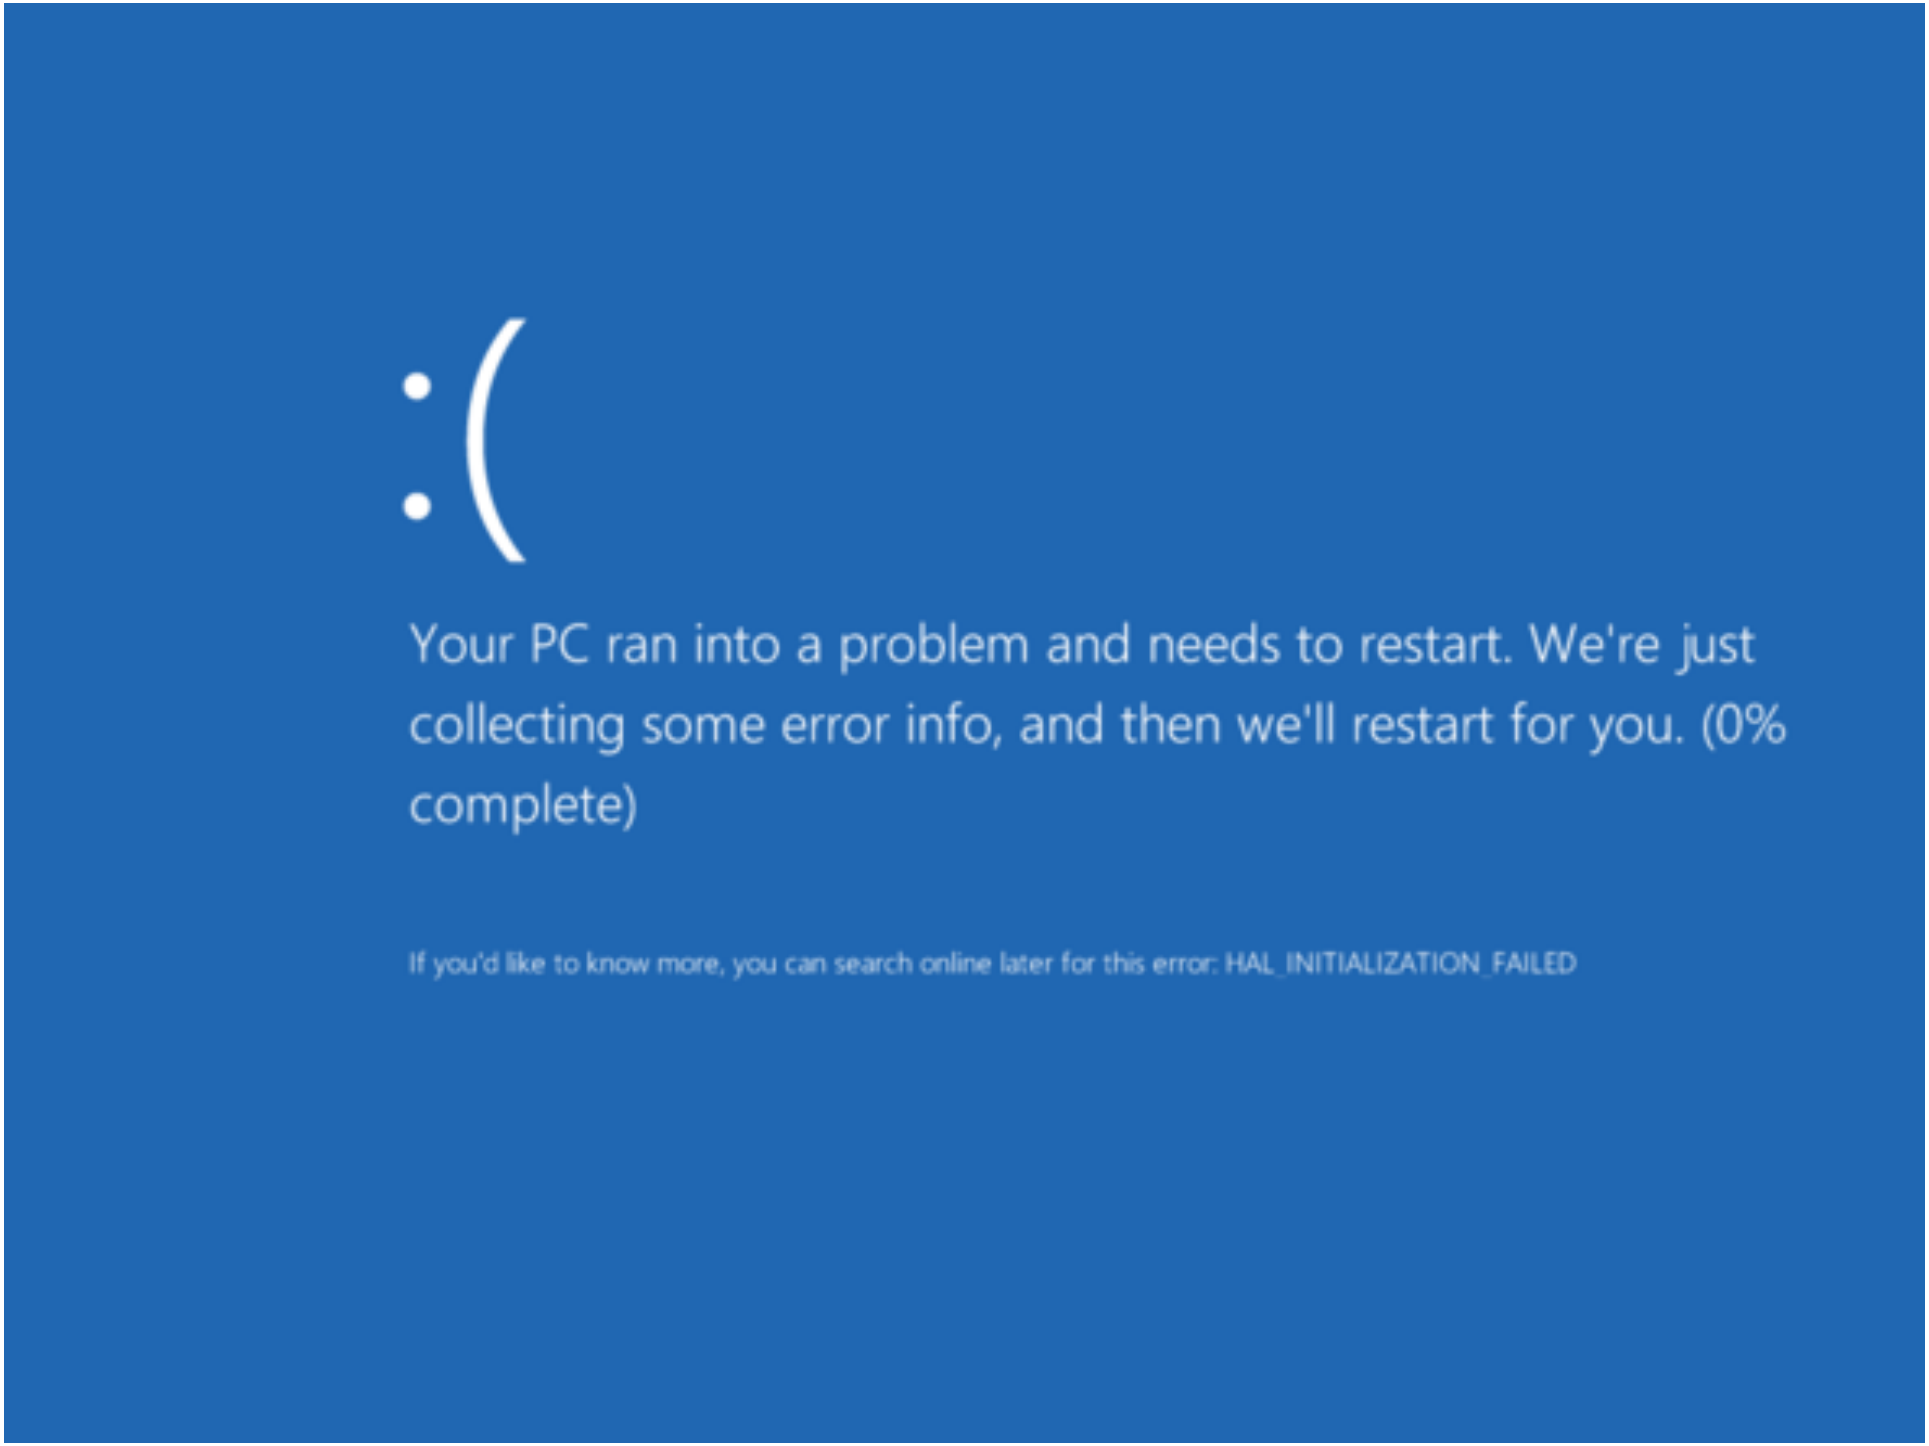
:(
Your PC ran into a problem and needs to restart. We're just collecting some error info, and then we'll restart for you. (0% complete)
If you'd like to know more, you can search online later for this error: HAL_INITIALIZATION_FAILED

San Francisco plane crash caused by pilot's inexperience with onboard computers

Boeing 777's auto-throttle did not maintain speed as expected

By Aaron Souppouris *on December 12, 2013 03:08 am* *Email* *@AaronIsSocial*

April 25, 1997 7:00 PM PDT

Router glitch cuts Net access

By CNET News.com Staff
Staff Writer, CNET News

Related Stories

Net blackout hits some regions

What started out as a router glitch at a small Internet service provider in Virginia today triggered a major outage in Internet access across the country, lasting more than two hours in some places.

What breaks in hardware?

COM1	
Component	%
Power supply	34.8
Memory	20.1
Hard drive	**18.1**
Case	11.4
Fan	8.0
CPU	2.0
SCSI Board	0.6
NIC Card	1.2
LV Power Board	0.6
CPU heatsink	0.6

10,000 machines

Pr(failure in 1 year) ~.3

Schroeder and Gibson, FAST 2008

Fault-Tolerant Systems from Unreliable Components

- TCP (retransmissions)
- DNS replication and caching
- MapReduce
- Resilient Overlay Networks

Availability in practice

- Carrier airlines (2012 FAA fact book)

Aviation Accident Rates by Type of Operation

Type of Operation	2007		2008		2009		2010		% Chg 10-09	
	Number	Rate	Number	Rate	Number	Rate	Number	Rate	Number	Rate
Large Air Carriers.....	28	.14	28	.15	30	.17	28	.16	-7%	-6%
Commuter.....	3	1.00	7	2.40	2	0.69	6	1.90	200%	175%
Air Taxi.....	62	1.50	58	1.80	47	1.60	31	1.05	-34%	-34%
General Aviation.....	1,652	6.90	1,566	6.90	1,474	7.20	1,435	6.87	-8%	-5%

Accident Rates are per 100,000 Flight Hours. Flight hours compiled by FAA.

Note: 2010 data are preliminary.
As of: 12/31/11

Source: NTSB
www.ntsb.gov/aviation/stats.htm

Other systems

- 911 Phone service (1993 NRIC report)
  - 29 minutes per line per year
  - ✓ 99.994%
- Standard phone service (various sources)
  - 53+ minutes per line per year
  - ✓ 99.99+%
- End-to-end Internet Availability
  - ✓ 95% - 99.6%


A19

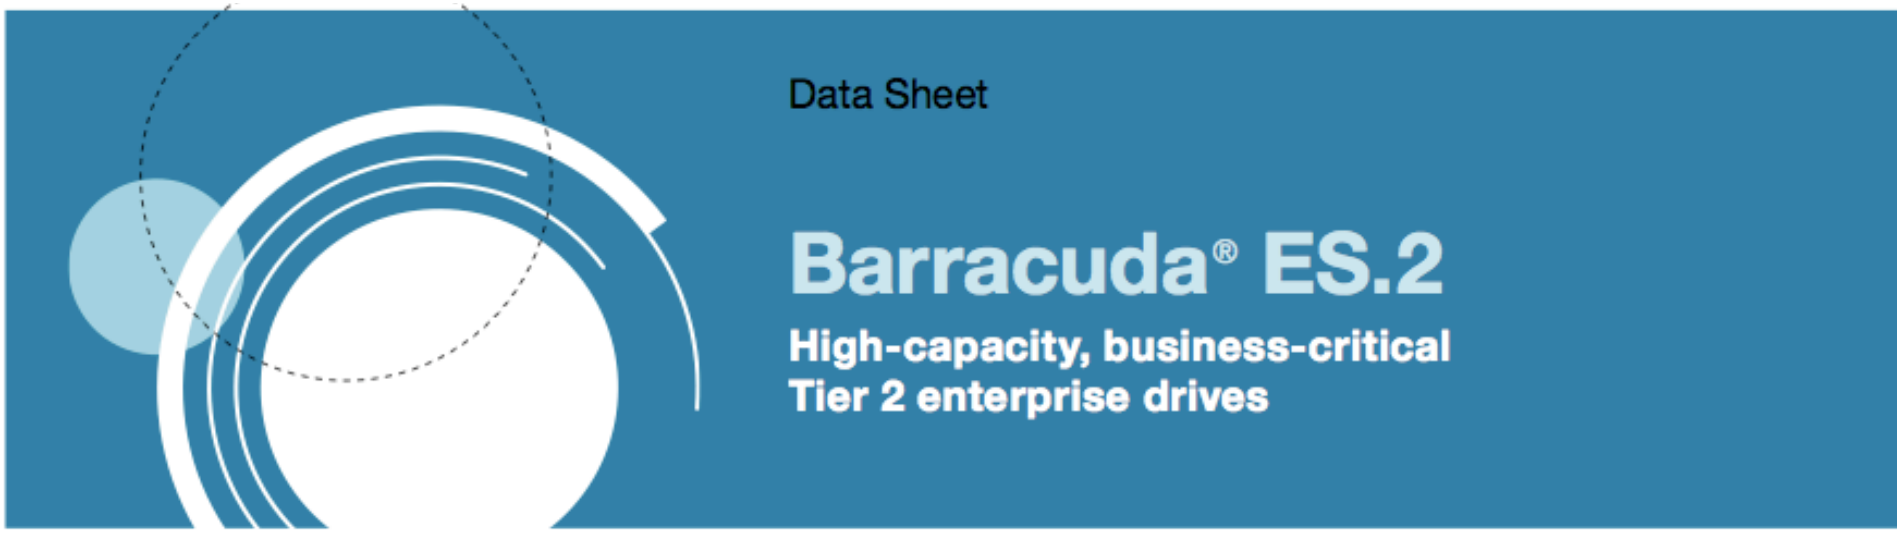

1 TB, 750 GB, 500 GB and 250 GB • 7200 RPM • SATA 3Gb/s, SATA 1.5Gb/s and SAS 3Gb/s

Reliability/Data Integrity	
Mean Time Between Failures (MTBF, hours)	1.2 million
Reliability Rating at Full 24x7 Operation (AFR)	0.73%
Nonrecoverable Read Errors per Bits Read	1 sector per 10E15
Error Control/Correction (ECC)	10 bit
Interface Ports	
SATA	Single
SAS	Dual

Party like it's 1879!

Cash register

First telephone
(Pres. Hayes)

Disk failure conditional probability distribution

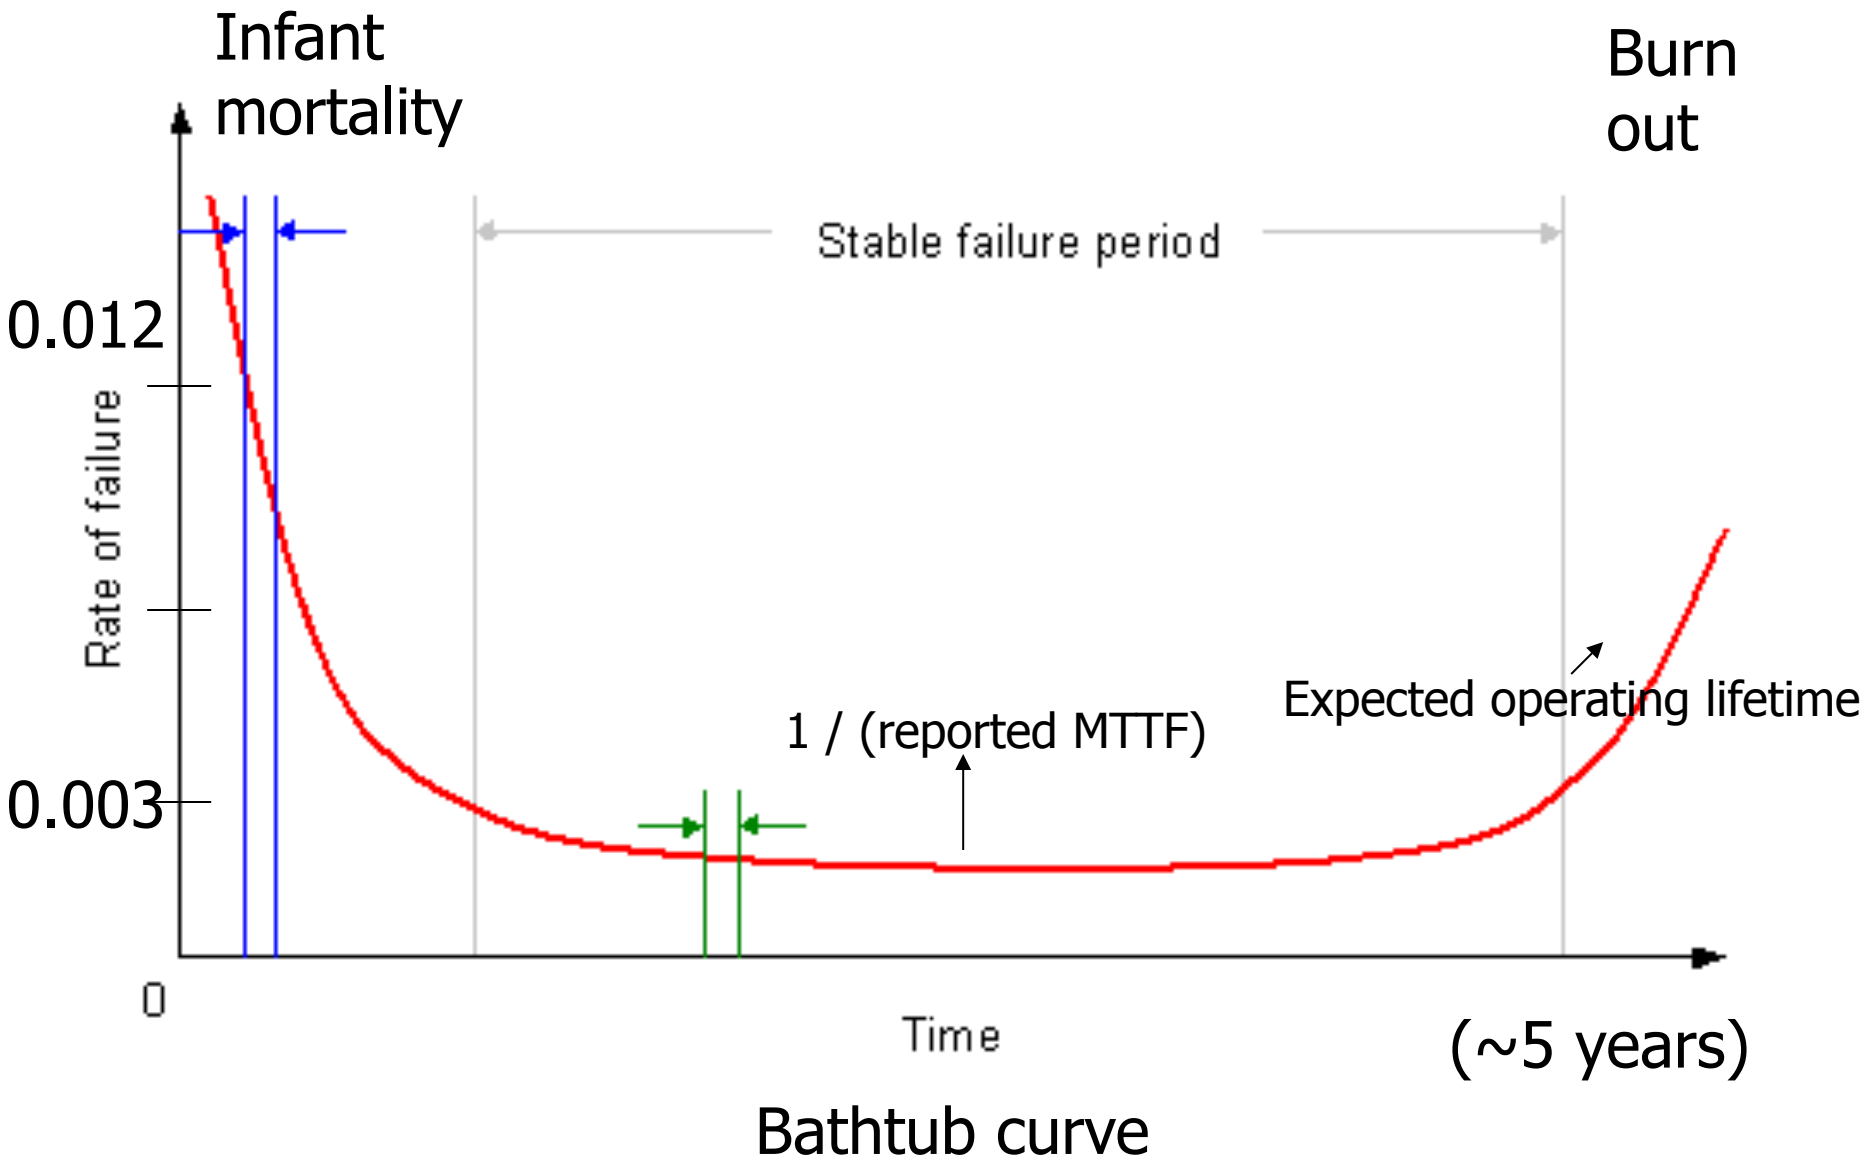

Bathtub curve

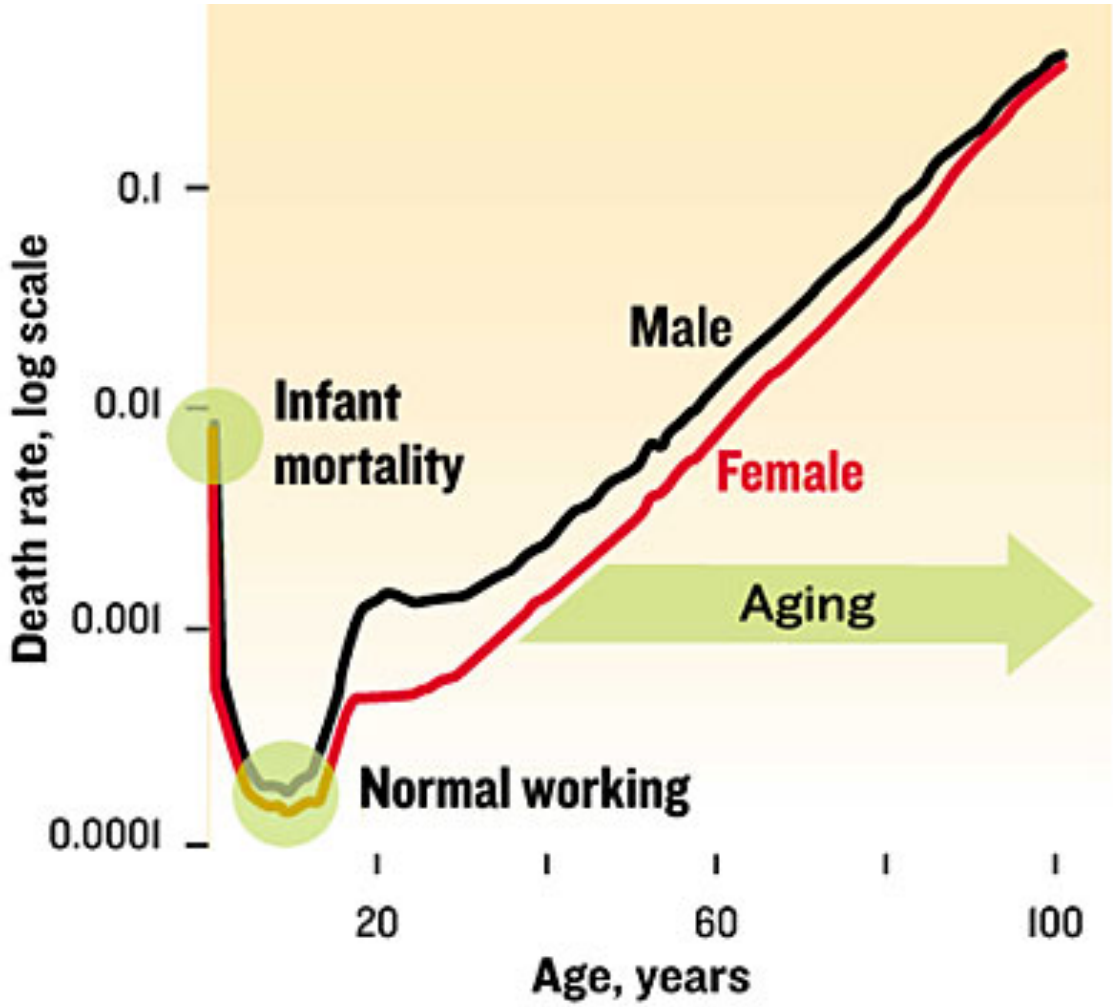

From: L. Gavrilov & N. Gavrilova, "Why We Fall Apart," IEEE Spectrum, Sep. 2004.
Data from http://www.mortality.org

Fail-fast disk

```
failfast_get (data, sector) {
        get (s, sector);
        if (checksum(s.data) == s.cksum) {
                data ← s.data;
                return OK;
        } else {
                return BAD;
        }
}
```

Careful disk

```
careful_get (data, sector) {
        r ← 0;
        while (r < 10) {
                r ← failfast_get (data, sector);
                if (r == OK) return OK;
                r++;
        }
        return BAD;
}
```

Replicated Disks
(RAID – recitation - more sophisticated)

```
write (sector, data):
    write(disk1, sector, data)
    write(disk2, sector, data)

read (data, sector):
    error = disk1.careful_get(data, sector)
    if error
        error = disk2.careful_get(data, sector)
        if error
            return error
    return data
```